

AD-A120 103 CARNEGIE-MELLON UNIV PITTSBURGH PA DEPT OF COMPUTER --ETC F/G 5/7
META-LANGUAGE UTTERANCES IN PURPOSIVE DISCOURSE.(U)
JUN 82 J G CARBONELL N00014-79-C-0661
UNCLASSIFIED CMU-CS-82-125 NL

[05]
AD A
120103

END
DATE
FILMED
DTIC

12

AD A120103

Meta-Language Utterances in Purposive Discourse

Jaime G. Carbonell

19 June 1982

DEPARTMENT
of
COMPUTER SCIENCE

DTIC FILE COPY



DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

DTIC
ELECTE
OCT 12 1982
B

Carnegie-Mellon University

82 10 12 037

Meta-Language Utterances in Purposive Discourse

Jaime G. Carbonell

19 June 1982

This research was sponsored in part by the Office of Naval Research (ONR) under grant number N0014-79-C-0661, and in part by the Defense Advanced Research Projects Agency (DOD), ARPA order number 3597, monitored by the Air Force Avionics Laboratory under contract number F-33815-81-K-1539. The views and conclusions in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Agency or the US Government.

References

Table of Contents

1. Introduction	1
1.1. Goal-determination inference	1
1.2. Social Role Constraints	3
1.3. Meta-language utterances	3
2. Interpreting Meta-Language Utterances	4
3. A Detailed Example	6
3.1. Local Ellipsis Resolution	6
3.2. Dialog-Level Ellipsis Resolution	7
3.3. Linguistic Referent Identification	7
3.4. Determining the Meta-Linguistic Action	7
3.5. Interaction with Goal Determination and Social Role Constraints	8
3.6. Interaction with Domain Knowledge	8
4. Meta-Language Utterance and Domain Knowledge	9
4.1. Implicit Linguistic Referent Identification	9
4.2. Resolving Indirect Speech Acts	9
4.3. The Need for Domain Knowledge	10
5. Towards a Taxonomy of Meta-Language Utterances	11
6. Concluding Remarks	13
7. References	14

Accession For	
NTIS GRAB	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
PER FORM 50	
By	
Distribution	
Availability	
Availability	
Dist	Special
A	

Meta-Language Utterances in Purposive Discourse

Jaime G. Carbonell

19 June 1982

Abstract

Meta-language utterances, i.e. utterances about other utterances, are ubiquitous in every type of human discourse, from unconstrained multi-party conversations to task-oriented man-machine dialogues. Several examples of meta-language communication are analyzed, identifying the complexities that a cognitively plausible discourse model should address. Meta-language utterances are investigated in conjunction with indirect speech acts and extralinguistic phenomena that impinge upon the discourse, such as social role constraints and goal determination inferences. It is argued that a comprehensive theory of discourse must necessarily address issues of meta-language communication, and linguistically-implicit communication conveyed through such means as inferred intent and cultural conventions.

1. Introduction

In order to formulate models of coherent discourse, a number of phenomena that arise in human dialogue must be analyzed from a computational standpoint. The pervasiveness of direct and indirect speech acts in unconstrained dialogue has made them a popular target for theoretical analysis [17, 18, 2], and computational modelling [13, 11], although a general computational model that interprets speech acts is yet to be perfected. However, this focus of attention on the part of researchers has tended to obscure the presence of other discourse phenomena, perhaps as ubiquitous and important as speech acts. Moreover, the *interplay* of various non-literal forms of communication may prove more critical in comprehending unconstrained human discourse than the analysis of any single discourse device.

Consider some of the more prevalent discourse phenomena, and the problems they pose for a computational model of a conversational participant, as illustrated in the examples below.

1.1. Goal-determination inference

An utterance may be interpreted differently depending on the inferred conversational goals of the speaker, other contextual knowledge being equal. Consider the following set of examples, in which the same utterance spoken in somewhat different contexts elicits radically different responses. These responses depend on the interpretation of the initial utterance, in which the attribution of goals to the

speaker plays a dominant role.

Passerby: *Do you know how to get to Elm Street?*

Person on the street corner:

You walk towards that tall building and Elm is your fifth or sixth left.

The passerby's question was quite naturally interpreted as an indirect speech act, since the information sought (and given) was not whether the knowledge of getting to Elm Street was present, but rather how to actually get there. Lest the mistaken impression be given that it is a simple matter to identify indirect speech acts computationally, consider the following variant to our example:

Passerby: *Do you know how to get to Elm Street?*

Person reading a street map and holding an envelope with an Elm street address on it:

No, I haven't found it; could you help me?

In the second example, the listener infers that the goal of the passerby is to render assistance, and therefore the initial utterance is interpreted as a direct query of the knowledge state of the listener, in order to know whether assistance is required. Hence, there is no indirect speech act in this example.

The issue becomes more complex upon considering the fact that the task of the interpreter involves more than extracting a binary decision on the nature of the speech act from goal expectations. The selection of *which* indirect speech act is meant often rests on contextual attribution of different goals to the speaker. Consider, for instance, the following contextual variant of our previous example:

Passerby: *Do you know how to get to Elm Street?*

Waiting cabbie: *Sure, hop in. How far up Elm are you going?*

In this example, the cabbie interpreted the goal of the passerby as wanting a ride to an Elm Street location. Making sure the cabbie knows the destination is merely instrumental to the inferred goal. The social relation between a cabbie and a (potential) customer is largely responsible for triggering the goal attribution. Thus, the passerby's utterance in this example is also interpreted as an indirect speech act, but a different one from the first example (i.e., be driven to the destination vs. knowing the location of the destination). In summary, three totally different speech acts are attributed to identical utterances as a function of different goals inferred from contextual information.²

²For additional discussion of goal-determination inferences in discourse comprehension see [1, 4, 6].

Example	Speech act
1) Original example	Indirect information request
2) Map reader	Direct information request
3) Cabbie example	Indirect action request

1.2. Social Role Constraints

The relative social roles of the discourse participants affects their interpretation of utterances as illustrated below:

Army General:	<i>I want a juicy Hamburger.</i>
Aide:	<i>Yes sir!</i>
Child:	<i>I want a juicy Hamburger.</i>
Mother:	<i>Not today, perhaps tomorrow for lunch.</i>
Prisoner 1:	<i>I want a juicy hamburger.</i>
Prisoner 2:	<i>Yeah, me too: All the food here tastes like cardboard.</i>

Clearly, the interpretation of the phrase "I want a juicy hamburger" differs in each example with no context present beyond the differing social roles of the participants, and their consequent potential for action. In the first example a direct order is inferred, in the second a request, and in the third only a general assertion of a (presumably unattainable) goal. Therefore, comprehending a dialogue rests critically on knowledge of social roles [4, 8]. Moreover, social role constraints provide part of the setting essential in making goal attributions, and therefore impinge (albeit indirectly) upon goal determination inferences discussed in the previous section. The following sections underscore the interaction between goal expectations, social role constraints, indirect speech acts, and meta-language utterance interpretation.

1.3. Meta-language utterances

Meta-language utterances consist of propositions about other utterances in the discourse, or about entire classes of potential utterances. In the course of a dialogue, direct references are often made to past utterances, such as "I meant to say edit, not delete", or "Forget I said that". These must be recognized as such and interpreted in the context of other conversational devices. For instance, the second utterance should not be interpreted literally as a command to erase an event from memory -- rather it is a request to disregard or undo the non-linguistic consequences of a past linguistic act.

I discussed goal-determination inference, social role constraints and meta-linguistic utterances as three examples of particularly significant, interacting open problems in discourse analysis, rather than an exhaustive enumeration of all such problems. Although a computationally effective model of

general discourse processing must address all issues that impinge upon the dialog [16], the bulk of my discussion focuses on problems that pertain directly to meta-language interpretation.³

2. Interpreting Meta-Language Utterances

On occasion, dialogue participants make statements about their past or future utterances, rather than addressing directly the topic under discussion. Consider, for instance, the following dialogue fragment between speakers A and B:

- A: *Please give me the Smithfield report*
 B: *Here it is, 'Smithfield 1982'*
 A: *I meant to ask for last year's report.*

The third utterance is, at face value, nothing more than a factual assertion. However, to understand it, one must find the referent linguistic act (i.e., the first utterance above), perform the necessary editing on it (first resolving the definite noun phrase "last year's report" to "last year's Smithfield report"), and then recognize and resolve the two indirect speech acts implied by the meta-language utterance, to wit: the speaker wants the consequences of the last action undone ("take back this year's Smithfield report") and a reinterpretation of the edited request ("Give me last year's Smithfield report.").

This example shows that a computationally effective model of meta-language utterances will involve coordinated solutions to at least the following problems:

- How to identify a meta-language utterance, as opposed to an assertion about a linguistic object playing no role in the discourse other than the presentation of information.
- How to recognize and confirm speech acts implicit in the meta-language utterance.
- How to detect and resolve ellipsis contained in the statement of the meta language utterance. (E.g., speaker A does not want "Last year's Smithfield report" to replace his entire first utterance; he wants that noun phrase to replace only the corresponding noun phrase "Smithfield report" in his earlier utterance -- a process akin to ellipsis resolution).
- How to determine the extension of a linguistic utterance onto the underlying world. (I.e., how to infer the non-linguistic consequences of a linguistic act that has been altered.)
- How to determine which consequences are no longer desired, and plan appropriate actions to undo them.

The above example also suggests that a discourse model incorporating a means of analyzing meta-

³ A forthcoming paper discusses issues of meta-language interpretation in limited, task-oriented domains, where simplified discourse structure and task constraints make the computational-modelling problem more tractable [3].

language utterances must be closely coupled with knowledge of actions, intent, and other dialogue phenomena (such as indirect speech acts), as well as more linguistic devices such as ellipsis and definite noun phrases.

Since meta-language pervades dialogue, its presence seldom rises to the awareness of participants. Consider, for instance, the following example (taken from a TV show):

- A: *Peter and John were already notified.*
B: *Did you say John already knows?*
A: *That's what I said; John knew about it before we did.*

The meta-language utterance "Did you say ..." in the dialogue fragment above presents some of the same problems as the earlier example: However, the question does not refer to the literal linguistic object just uttered; it refers to one of the major consequences of the of the underlying action (i.e., augmenting John's knowledge state). Speaker A did not literally say "John already knows...", but since this is a direct inference from his utterance, it is treated as equivalent to what he actually said at the dialog-level analysis. Hence, the referent of the meta-utterance cannot be established without appeal to an underlying model of actions in the established domain, as well as notions of topic and focus [6, 7, 19].

On occasion, several meta-utterances intermesh, as in the dialogue fragment below:

- A: *John said you were going to dance with Mary today.*
B: *He couldn't have said that; he knows I'm going elsewhere.*
A: *I think he said you were taking Mary to the dance.*
B: *He meant I was driving her there on my way to card game.*

This example is particularly complex, in that every utterance is a meta-utterance, all referring to a unique linguistic object which itself is not part of the dialogue proper. Once again, ellipsis resolution, inferences based on an underlying action and intentional model, and definite noun phrase resolution are required components of an analysis. However, the coherence of all four meta-utterances into a meaningful dialogue fragment defies attempts at localized analysis. This example is meant to illustrate the complexities that arise in meta-language use in unconstrained discourse.

Meta-language may also affect the interpretation of future discourse. In this case, the referent "linguistic object" is typically an entire class of utterances fitting a given description. Examples of this form of *indefinite meta-language utterances* are presented with minimal commentary among the examples below:

"When I say I want something, I mean for you to get it for me immediately!" establishing explicitly what may otherwise be an indirect speech act

"Do not take what I say literally." clearly affecting how the other participant interprets future utterances in the dialogue

"Can't you understand what 'Do you have the time' means?"

"Did you say the messages 'to' John or 'from' John?"

"I should not have typed 'delete'." Note the different responses called for if the listener is a natural language interface to an software system, a teacher helping the person do better next time, or a casual friend trying to ease the anxiety of the speaker.

Having discussed the range of meta-language utterances in discourse, let us look more closely at the processing the must underlie meta-language interpretation, including linguistic and extra-linguistic knowledge that must be brought to bear. As discussed earlier, a general process model must incorporate many interacting discourse devices, and is therefore beyond the current stage of our investigation. Hence, I discuss processing issues in meta-language interpretation by analyzing sample discourse fragments from a domain where social roles and implicit goals are simple and easily discerned.

3. A Detailed Example

Consider the following example of a meta-language utterance (MLU) in the context of a natural language interface to a simulated operating system:

USER: *Copy the file GAUSS.FOR from the system library to my directory.*

SYSTEM: *Non existent file: "GAUSS.FOR"*

USER: *Oops! I meant GAUSS.DAT.*

If the system were to respond with a remark like "Fact recorded: you meant GAUSS.DAT", the user would be less than satisfied. Why? Clearly the user meant to convey information relevant to a reinterpretation of his previous utterance. The issue involved here is more complex than the mere presence of an indirect speech act. To wit, the following processes must be invoked.

3.1. Local Ellipsis Resolution

First, the local semantic ellipsis must be resolved. The semantic object of the verb "to mean" requires a conceptualization (i.e., the expression of an action or state change). Thus the MLU may be expanded to: "I meant *to type* GAUSS.DAT"

3.2. Dialog-Level Ellipsis Resolution

There is yet another ellipsis, this one at the *dialog level*. The expanded MLU may be a semantically complete unit at the sentential level, but not a complete unit with respect to the dialog or with respect to meta-language evaluation. Let me elaborate: If one were to evaluate the local intent expressed in "I meant to type GAUSS.DAT", the resultant action would be actually typing to the system "GAUSS.DAT". However, "GAUSS.DAT" would then be interpreted as an isolated and incomplete utterance. Additionally, the context of the dialog expects a possible resolution to the problem caused by the non-existent file, and the "oops! I meant..." cues that an error recovery process may be in effect. Both of these concerns suggest that the utterance must be understood in the context of the dialog.⁴ These cues identify the utterance as a potential MLU.

3.3. Linguistic Referent Identification

Once a potential MLU utterance is identified, its linguistic referent must be found. Finding the referent utterance (or set of utterances) confirms the presence of a MLU. In our example, the linguistic referent is the user's previous utterance. This allows one to resolve the discourse level ellipsis, and the user's second utterance is transformed in essence to: "I meant to type GAUSS.DAT *instead of something I typed as part of my previous utterance*". The fact that the previous utterance resulted in an error, coupled with the error-correction cues mentioned earlier, facilitate the identification of the linguistic referent.

3.4. Determining the Meta-Linguistic Action

Once the referent is found, the appropriate meta-linguistic action must be performed. In the present example, that meta-linguistic action is a lexical substitution preserving syntactic and semantic integrity. "GAUSS.DAT" replaces a linguistic object of like syntactic and semantic type, namely the previous file specification "GAUSS.FOR". Note that the MLU would have been ambiguous if there had been more than one semantically and syntactically legal substitution, and meaningless if there had been no legal substitution. This substitution process results in the the second utterance being fully interpreted: "I meant to type GAUSS.DAT *instead of GAUSS.FOR* in my previous utterance." Evaluating the resolved MLU on its linguistic referent yields a new modified utterance: "Copy the file

⁴The reader should note that I am proposing a departure from the more conventional strategy of understanding utterances fully at the sentential level, and only then interpreting their discourse import. Rather, the actual meaning representation for MLUs can only be captured by relating them to their dialog referents. This differs from the process of creating a local structure that is then embedded in a larger structure (as in simple ellipsis resolution), or identifying anaphoric or definite noun phrase referents, where the parsed structure is left intact. In these cases, a local structure is created, irrespective of dialog context. That structure remains intact and accrues information when processed at the dialog level. We are arguing that for MLU resolution, one *cannot* build a local structure without *first* identifying dialog referents. Thus, a more integrated process model is required to handle such complex discourse-level phenomena.

GAUSS.DAT from the system library to my directory."

3.5. Interaction with Goal Determination and Social Role Constraints

Now that the meta-language has been processed, one is still faced with the discourse imperatives: namely, the new utterance must be interpreted in light of the goals of the user, the master-slave social relation between the user and the system, and the ensuing indirect speech act. If the utterance had occurred in the context of a user communicating with a text-editing system, "I mean to type..." could simply result in correcting the previous text (after the MLU identification process discussed above). In other words, the goal of the user in a text editing system is to create or alter text. Once such an alteration is specified by the MLU referent identification process, it matches a known goal and the relation between user and system is invoked -- namely, the role of the latter is to implement the goals of the former. With both goal and method specified, the system need only correct the previous text and be satisfied.

However, the present goal of the user is to avail himself of the functionality provided by an operating system; thus it should be inferred that the user means to reexecute modified commands. Thus, the act of altering the text of his previous utterance does not in itself a goal, and the goal determination process must infer that since user goals in an operating system interface are achieved by executing commands, the altered text should be interpreted as a new command to be executed. Having inferred the indirect speech act that the user is requesting a command reexecution, the system, at the beck and call of its master,⁵ will perform the necessary actions. In general, after MLU identification, a speaker's MLU utterance is interpreted as the indirect speech act appropriate in the current goal-expectation setting, consistent with the social roles of the discourse participants.

3.6. Interaction with Domain Knowledge

The process of reexecuting the modified copy command occurred in the underlying (non-linguistic) system, which was treated as a proverbial black box by the discourse analysis processes. In the present example, all the difficult problems occurred at the dialog level, necessitating no interaction with the underlying system other than transmitting the resultant reinterpreted command to the non-linguistic system. However, this is not generally the case, as we shall in the following variant of our example.

⁵Recall the simplified social role constraints of the present natural language interface scenario, wherein the role of the system is that of a slave to the user. It would indeed be a strange and unwelcomed interface that decided to pursue goals other than those of the user and responded with "copy that file yourself!"

4. Meta-Language Utterance and Domain Knowledge

In the example above, MLU interpretation interacts with goal settings, with the implicit (and rather one-sided) social relation, and with indirect speech act interpretation. Hence, to comprehend an apparently simple dialog fragment, easily understood by humans, requires analysis of dialog devices well beyond the scope of presently implemented computer models -- and somewhat beyond the boundaries of present theoretical analysis. To illustrate this point more emphatically consider a variant of the previous example, in which knowledge of the underlying domain itself becomes crucially important.

USER: *Copy the file GAUSS.FOR from the system library to my directory.*

SYSTEM: *File GAUSS.FOR copied to your directory.*

USER: *Oops! I meant GAUSS.DAT.*

This example is identical to the previous one, except that no error occurred in copying the file originally specified. Does the analysis differ? All the issues of ellipsis resolution, MLU recognition, linguistic referent identification, meta-linguistic action interpretation, and goal determination inference discussed above apply in the new example. In addition, a second indirect speech act is present, and more substantial knowledge of the underlying domain is required. More explicitly, the additional complications include the following:

4.1. Implicit Linguistic Referent Identification

The linguistic referent identification is made more difficult. Whereas in the previous example the system could infer the meta-linguistic action to be an error-correcting process and hence focus on the culprit utterance, there is no explicit error in any of the previous utterances of the new dialog fragment. The meta-linguistic action is still one of error-correction by means of lexical substitution, but now the error exists only in the mind of the user, and the system must find the probable linguistic referent by other means. Here some of the focus tracking and definite noun phrase literature [14, 7, 19] could come to the rescue, if suitably extended to apply to meta-linguistic constructions.

4.2. Resolving Indirect Speech Acts

Once the linguistic referent of the MLU is found, the meta-linguistic action of lexical substitution is carried out and other dialog devices including goal-directed inference and social role constraints come into play, as in the previous example. However, in the present case there are *two distinct indirect speech acts* to be identified. The first is that the user, by signaling his error and relying on the system to infer and carry out its goals, intends the consequences of his previous actions to be undone. To wit, the user means for the system to delete the copy of the unwanted file GAUSS.FOR.

The process of identifying this indirect speech act is a general one of appealing to the converse of the principle of rationality [5, 21, 12]. Simply stated, that principle asserts that all actions of rational entities are in direct service to their goals. Therefore, its converse states that actions not performed in service of goals are purely accidental, and their consequences are, in general, not desired. Hence the first indirect speech act may be resolved as an implicit request to take back the consequences of the undesired action that resulted from the previous command. The second indirect speech act is, as before, carrying out the reinterpreted, modified command. Thus, the indirect speech acts must be recognized by appealing to knowledge to the user's goals, general principles of rational behavior, and knowledge of the social relation between the user and system.

4.3. The Need for Domain Knowledge

In order to instantiate the true meaning of the speech act indirectly requesting that the consequences of the previous act be undone, knowledge of the domain is required. The system must know that:

- The action of **COPYING** entails lasting consequences of potential significance.
- The primary consequence consists of creating a file whose contents are those specified by the *object slot* of the **COPY** case frame, and whose location is specified by the *destination slot* of the **COPY** case frame.
- The way to eliminate an unwanted file is to apply the **DELETE** command with its *object slot* set to the offending file name and the *source slot* set to the location of the unwanted file. Thus, domain knowledge, together with at least a rudimentary planning capability in order to bring that knowledge to bear, are needed to interpret fully the discourse import of the MLU, and thereby generate appropriate action.

The component of the system that interacts with the underlying domain cannot be divorced from the dialog analysis process -- although one can certainly formulate a theory of discourse phenomena independent of any application domain. The distinction I advocate is that the discourse model must have the ability to access domain knowledge when needed, but must not *incorporate* internally any such knowledge. Thus, a *process model* of discourse incorporates only techniques for tracking context and focus, and for recognizing, interpreting, and responding to discourse devices, plus mechanisms to interact with, and utilize information from external modules, such as goal analysis, planning, specific domain knowledge, knowledge of a particular language, and knowledge of social interactions.

5. Towards a Taxonomy of Meta-Language Utterances

A comparative analysis of a large set of meta-language utterances (MLUs), including those presented above, suggests that they may be usefully classified according to three (not truly orthogonal) primary criteria:

- The type of linguistic object that the utterance operates on. Some MLUs refer to a specific past utterance, others define an utterance for future reference, and yet others refer to entire classes of utterances.
- The role played by the MLU in the context of the discourse. The roles may be quite varied, as discussed below.
- The underlying action(s) expected on the part of the listener of the MLU.

It appears that classifying each MLU along these three dimensions captures the bulk of the variability observed for the set of examples I have analyzed. At this point in the investigation, I am not certain how to exploit this linguistic taxonomy in deriving a cognitively plausible model of human MLU interpretation. The outline below provides a rather detailed taxonomy of MLUs organized around the three dominant dimensions.

1. Linguistic object referent.

a. Definite referent MLU

- i. The referent is within the scope of the discourse.
This is the most common type of MLU, as illustrated in most of the examples
- ii. The referent is outside the scope of the discourse
The complex "dance" conversational fragment exemplifies a dialogue where the MLUs refer to a linguistic object not part of the dialogue proper

b. Indefinite referent MLU

- i. The referent is a *class* of linguistic objects
E.g., "When the boss asks for something, you get it for him immediately."
- ii. The referent is introduced by the MLU
This is analogous to indefinite noun phrases. E.g., "Start the engine when I say 'go'."

2. Role of the MLU in the discourse

a. Structuring role

- i. The linguistic object of the MLU becomes topic of the dialogue as in the "dance" example

b. Repair past errors or misunderstandings

i. Modifications of previous utterances.

1. Correction of interpretation errors

E.g., "When I said go, I meant for you to start the engine, not leave the room."

2. Explicit changes to the linguistic object itself with implicit semantic reinterpretation.

E.g., "Oops! I meant to type BAZ.FOO."

ii. Deletions of (the semantic consequences of) previous utterances.

E.g., "Forget I said that" or "Nevermind what I said"

iii. Insertions of utterances into the past discourse.

E.g., "Oh, I should also have said that the problem assumes adiabatic expansion of gases."

c. Change future interpretation strategies by the listener

i. Tell the listener to employ a particular mode of interpreting following utterances.

E.g., "Metaphorically speaking ...", "In gutter talk I would say..."

ii. Speaker-specific interpretation

E.g., "Never take John literally", or "Mary speaks in riddles"

d. Define new language constructs

E.g., "JFK is an abbreviation for John F Kennedy", "A grass snake is a snake that lives in the grass, not one that eats grass.", or "When the boss says he wants something, you jump and get it for him!"

3. Type of action taken in response to the MLU

a. A linguistic Act such as redirecting the discourse

b. An overt act such as handing over the appropriate "Smithfield report" in an earlier example

c. An internal change such as remembering that 'go' means to start the engine in the context of the present discourse

d. A social act such as reaffirming social roles, e.g. "I said I wanted a **juicy** hamburger, Corporal!"

e. Any combination of the above.

6. Concluding Remarks

I have argued that Meta-Language Utterances pervade unconstrained discourse, and their comprehension requires a unified model of discourse, incorporating aspects of ellipsis resolution, definite noun phrase resolution, direct and indirect speech acts, goal determination inference, social role constraints, and interaction with topic-selection and domain semantics processes. Given the complexity of such a comprehensive model, it behooves us to continue theoretical investigation of all relevant discourse phenomena. Each (apparently independent) discourse device impinges upon other constituents and thereby provides another piece of the global puzzle. I envision the role of MLUs in discourse analysis to be threefold:

1. Continue to investigate their function, structure and linguistic implications in greater depth than the current study
2. Use MLU interpretation as a method of testing the completeness and fine-structure detail⁶ of proposed theories of discourse. Given their complexity and their intimate association with other discourse devices, MLUs play a rather central role in discourse. Therefore, any theory that fully addresses MLU comprehension has made long strides towards comprehensive coverage of discourse phenomena. Conversely, if that theory does not provide for the level of analysis required for MLU comprehension it lacks the requisite commitment to fine-structure detail necessary for either completeness or falsifiability.
3. In very narrow application domains, where the goals of the participants, their social interactions, and the topic of discussion are significantly constrained, I propose to start developing a computationally effective model of MLU interpretation. These situations do not truly reflect unconstrained human discourse, but occur often in natural language-based man-machine interfaces. The computational investigation of meta-language interpretation in constrained domains may shed light onto the possible cognitive processes whereby humans comprehend MLUs with little apparent difficulty. Moreover, I have found that MLUs account for a significant fraction of utterances that present task-oriented natural language interpreters are unable to process correctly. Therefore, practical benefits may be reaped from a computational analysis, as well as insights of a more theoretical nature.

In addition to the forms of inter-sentential meta-language utterances discussed in this paper, the use of intra-sentential meta-linguistic devices deserves study. Instances of intra-sentential use of meta-language include the well known "respectively" examples, and sentences such as "John, Bill and Mary will visit us *in that order*." [9, 15]. Intra-sentential meta-language affects sentential-level analysis, inter-sentential meta-language affects discourse-level analysis. Whereas inter-sentential meta-language pervades dialog, it is more seldom found in written text, where intra-sentential meta-

⁶Sloman [20] argues that fine structure is a useful measure of theoretical adequacy absent from most analyses in the Philosophy of Science, such as Kuhn's treatise [10]

language occurs with undiminished frequency. Therefore, inter-sentential meta-language utterances are truly a discourse phenomenon and can be meaningfully studied only in the context of a more comprehensive model of human discourse.

Acknowledgement

I would like to thank Phil Hayes who provided some helpful insights and helped make earlier drafts of this paper more coherent.

7. References

1. Allen, J. F. and Perrault, C. R., "Analyzing Intention in Utterances," *Artificial Intelligence*, Vol. 15, No. 3, 1980, pp. 143-178.
2. Austin, J. L., *How to Do Things with Words*, Oxford University Press, 1962.
3. Carbonell, J. G. and Hayes, P. H., "Metalanguage in Discourse," Tech. report, Dept. of Computer Science, Carnegie-Mellon University, 1982, (forthcoming).
4. Carbonell, J. G., *Subjective Understanding: Computer Models of Belief Systems*, Ann Arbor, MI: UMI research press, 1981.
5. Carbonell, J. G., "Where do Goals Come From?," *Proceedings of the Fourth Annual Meeting of the Cognitive Science Society*, 1982, Ann Arbor, MI.
6. Frederking, R., "A Rule-Based Conversation Participant," *Proceedings of the 19th Meeting of the Association for Computational Linguistics*, 1981.
7. Grosz, B. J., *The Representation and Use of Focus in Dialogue Understanding*, PhD dissertation, University of California at Berkeley, 1977, SRI Tech. Note 151.
8. Grosz, B. J., "Utterance and Objective: Issues in Natural Language Communication," *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, 1979, pp. 1067-1076.
9. Joshi, A. K., "Use (or Abuse) of Metalinguistic Devices," 1972. Unpublished Manuscript.
10. Kuhn, T. S., *The Essential Tension: Selected Studies in Scientific Tradition and Change*, U. of Chicago Press, 1977.
11. Levin, J. A., and Moore, J. A., "Dialogue Games: Meta-Communication Structures for Natural Language Understanding," *Cognitive Science*, Vol. 1, No. 4, 1977, pp. 395-420.
12. Newell, A., "The Knowledge Level," Tech. report, Dept. of Computer Science, Carnegie-Mellon University, 1981, CMU-CS-81-131.
13. Perrault, C. R., Allen, J. F. and Cohen, P. R., "Speech Acts as a Basis for Understanding Dialog Coherence," *Proceedings of the Second Conference on Theoretical Issues in Natural Language Processing*, 1978.

14. Ross, J. R., "On the Cyclic Nature of English Pronominalization," in *To Honor Roman Jakobson*, The Hague: Mouton & Co., 1967.
15. Ross, J. R., "Metaanaphora," *Linguistic Inquiry*, 1970 .
16. Schank, R. C. and Carbonell, J. G., "Re The Gettysburgh Address: Representing Social and Political Acts," in *Associative Networks*, Findler, ed., New York: Academic Press, 1979.
17. Searle, J. R., *Speech Acts*, Cambridge University Press, 1969.
18. Searle, J. R., "Indirect Speech Acts," in *Syntax and Semantics, Volume 3: Speech Acts*, P. Cole and J. L. Morgan, eds., New York: Academic Press, 1975.
19. Sidner, C. L., *Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse*, PhD dissertation, MIT, 1979, AI-TR 537.
20. Sloman, A., *The Computer Revolution in Philosophy: Philosophy, Science And Models of the Mind*, Harvester Press, 1978.
21. Wilensky, R., *Planning and Understanding*, Addison Wesley, Reading, MA, 1983.

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER CMU-CS-82-125	2. GOVT ACCESSION NO. AD-A120103	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) META-LANGUAGE UTTERANCES IN PURPOSIVE DISCOURSE		5. TYPE OF REPORT & PERIOD COVERED Interim
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) JAIME G. CARBONELL		8. CONTRACT OR GRANT NUMBER(s) O N0014-79-C-0661
9. PERFORMING ORGANIZATION NAME AND ADDRESS Carnegie-Mellon University Computer Science Department Pittsburgh, PA. 15213		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS		12. REPORT DATE June 19 1982
		13. NUMBER OF PAGES 18
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report)		
<div style="border: 1px solid black; padding: 5px; text-align: center;"> DISTRIBUTION STATEMENT A Approved for public release; Distribution Unlimited </div>		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
Approved for public release; distribution unlimited		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		
<p>Meta-language utterances, i.e. utterances about other utterances, are ubiquitous in every type of human discourse, from unconstrained multi-party conversations to task-oriented man-machine dialogues. Several examples of meta-language communication are analyzed, identifying the complexities that a cognitively plausible discourse model should address. Meta-language utterances are investigated in conjunction with indirect speech acts and extralinguistic phenomena that impinge upon the discourse, such as social role constraints and goal determination inferences. It is argued</p>		

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

that a comprehensive theory of discourse must necessarily address issues of meta-language communication, and linguistically-implicit communication conveyed through such means as inferred intent and cultural conventions.

S. N 0102- LF- 014- 6601

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

END

DATE
FILMED

11-82

DTIC